

On a Unitary Semantical Analysis for Definite and Indefinite Descriptions

Peter Ludlow and Gabriel Segal

1. Introduction

Philosophers of language (and semanticists) do not agree on much, but few have felt reason to doubt that there are at least two kinds of descriptions in natural language: definite descriptions (e.g. of the form ‘the F ’), used in sentences which say that there is a unique satisfier of F , and indefinite descriptions (e.g. of the form ‘an F ’), used in sentences which claim only that something or other satisfies F . However, if analytic philosophy had spawned in languages other than English and German, this assumption might not appear so innocent. In many languages—even many Indo-European languages—we do not find these two kinds of descriptions, at least not on the surface. In Russian, for example, neither definite nor indefinite articles appear on the surface: ‘the man’ and ‘a man’ would both be expressed by the Russian equivalent of ‘man’.

If we suppose that natural languages do not differ from each other in their logical form,¹ then there are a limited number of possible paths we can take to explain the discrepancy between English and these other languages. The first path is simply to say that, for example, Russian *has* both definite and indefinite descriptions, but that they do not appear on the surface; they are simply unpronounced lexical items. In other words, in Russian and similar languages, there is more semantic structure than meets the eye.

We are indebted to Delia Graff and Stephen Neale for discussion, to Kent Bach, Emma Borg, Ray Elugardo, Kasia Jaszczolt, Ruth Kempson, Mike Nelson, and Jason Stanley for comments on an earlier version of this chapter, to Anne Bezuidenhout and an anonymous reviewer for comments on the penultimate draft, and to Zoltán Szabó for comments on multiple drafts.

¹ This assumption could be challenged, of course, but work in generative linguistics over the last three decades has convinced us that differences between natural languages are often superficial, and that at the level of description we are interested in those differences are vanishingly small. The thesis is empirical, however, and may yet turn out to be false, but it has been fruitful to date and we accordingly adopt it here.

378 *IV. Definites and Indefinites in Semantic Theory*

The second path is to say that languages like English that *appear* to have definite and indefinite determiners do not actually have two determiners in the sense of two logical operators with distinct semantical contributions. Rather, we can say that English has two surface grammatical elements ('a' and 'the') with the same literal meaning. As we noted above, few have taken this option, but it has been on the table for some time. Earlier attempts to defend the thesis include Kempson (1975), Breheny (1999), Szabó (2000), and Zvolensky (1997), and the idea is at least entertained in Heim (1982) and Kamp and Reyle (1993). In this chapter, we revisit the position and enlist a number of new arguments in defense of it.

Before we get into the arguments for the proposal, it is first necessary to see that there is a great deal of intuitive appeal to the idea (as strange as it may seem at first blush). Accordingly, we will begin with a discussion motivating the general position, and a possible technical execution of the idea. We will then turn to some of the arguments for the position (in section 3) and will conclude (section 4) with some general observations about why the idea initially seemed so strange.

2. The Intuitive Appeal of the Basic Thesis

The basic idea under consideration is that for both 'An *F* is *G*' and 'The *F* is *G*' the literal semantics is exhausted by the interpretation that $[\exists x: Fx][Gx]$ receives in a standard truth-conditional semantics. In other words, in English there is *less* semantic structure than meets the eye.²

It has been widely noted that there are many cases in which we use definite determiners without expressing uniqueness. Consider the following examples:³

- (a) John went to the dentist.
- (b) John was hit in the eye.
- (c) The boy scribbled on the living-room wall. (Du Bois 1987)
- (d) Towards evening came to the bank of a river. (Christopherson 1939)
- (e) Let's go to the pub. (British English)
- (f) John drove into the ditch and had to go to the hospital.

² A third possibility, consistent with what we say here, is that there is no genuine Russellian determiner at all (either definite or indefinite) but rather that the surface determiners are, strictly speaking, free variables which might be unselectively bound by discourse elements, and which, when unbound, would be interpreted like an indefinite description. Heim (1982) and Kamp (1981) have proposed this analysis for the determiner 'a', but so far as we know, a similar proposal has not been made for the determiner 'the' (see Kamp and Reyle 1993: ch. 3) for discussion of puzzles surrounding the treatment of 'the' in Discourse Representation Theory, DRT). The point of this chapter for DRT would be that the analysis can be extended likewise to the determiner 'the'. We pass over that possibility here because we do not want to be deflected by discussions of the relative merits of DRT.

³ We are indebted to Anne Bezuidenhout for bringing some of these examples to our attention and to Abbott (1999) for tracking down their sources in the literature.

- (g) Take the elevator to the sixth floor. (Birner and Ward 1994)
 (h) No problem, I'll get the maid to do it. (Epstein 2000)

Our view is that these cases are not aberrations, but in fact more closely reflect the semantics of so-called definite determiners. What we intend to show is that in those cases where there is apparent definite force or uniqueness, pragmatic considerations are in play.

The basic argument is straightforward. Just as Kripke (1977) argued that we need not posit an ambiguity in 'the *F*' to account for the referential and attributive uses of descriptions, and just as Ludlow and Neale (1991) argued that we need not posit a semantical ambiguity to account for the referential, specific, definite, and quantificational uses of 'an *F*', one can likewise argue that one need not draw on a semantic distinction between 'the *F*' and 'an *F*' to account for the different uses to which these descriptions are applied. Specifically, 'the *F*' and 'an *F*' literally make the same contribution to truth conditions: they both have the semantics that we would ordinarily represent as $(\exists x)Fx$ in first order logic. The *definite* use of both 'the *F*' and 'an *F*' is derived by standard Gricean pragmatic resources.

3. Speaker's Reference/Semantic Reference and Beyond

Donnellan (1966) pointed out that definite descriptions can be used in (at least) two different ways. On a so-called *attributive* use, a sentence of the form 'The *F* is *G*' is used to express a proposition equivalent to 'Whatever is uniquely *F* is *G*'. For example, on seeing Smith's badly mutilated corpse, Robinson might say 'The murderer of Smith is insane' thereby expressing the idea that some unique individual murdered Smith and whoever that individual is is insane.

On a *referential* use, a sentence of the form 'The *F* is *G*' is used to pick out a specific individual, *x*, and say of *x* that it is *G*. For example, suppose Jones is on trial for Smith's murder and is ranting and raving in the dock. Brown says 'The murderer of Smith is insane', thereby expressing the idea that Jones is insane. The proposition thus conveyed is true whether or not Jones murdered Smith. Donnellan suggested that Russell's quantificational account of definite descriptions might capture attributive uses, but that it does not seem right for referential uses.

Kripke (1977) responded by arguing that a univocal, quantificational account of definite descriptions could account for both referential and attributive uses. The distinction between the two uses would thus be a purely pragmatic matter. In essence, Kripke's argument goes like this. Assume that 'the' has a univocal, quantificational semantics. Then speech acts and acts of communication would still proceed in the way Donnellan says they do. Suppose, by hypothesis, that at the trial when Brown says 'The murderer of Smith insane' the literal meaning of the words is that exactly one person, *x*, is such that *x* murdered Smith and *x* is insane. Brown

380 *IV. Definites and Indefinites in Semantic Theory*

would still, in that context, succeed in communicating the proposition that Jones is insane.

How does this work? The basic idea is simple. Brown speaks while Jones rants and raves in the dock, or shortly after. Following Grice's (1975) maxim of quality we may suppose that Brown is attempting to say something true. He must, therefore, believe that exactly one person, *x*, murdered Smith and *x* is insane. He must have grounds for this belief. Now either these grounds are specific, based on beliefs about some particular individual believed to be the murderer, or they are not, in which case they would be based on general features of the case (such as the appalling condition of Smith's body). But there is no evidence that Brown has any general grounds for the belief and good evidence that he has specific grounds—for someone accused of Smith's murder is highly salient and acting strangely.

Moreover, in the particular conversational context, an expression of a belief based on purely general grounds would not be relevant: general features of the case are not under discussion at the time. The audience may thus infer that Brown would not have said what he did unless he believed that Jones was insane. Moreover Brown knows all of the above and knows that his audience knows it, and so on. Therefore, Brown has communicated the proposition that Jones is insane.

Kripke's argument is that the univocal semantics plus independently motivated pragmatics suffices to account for the data. There is therefore no need to posit an ambiguity in 'the'. Doing so would be redundant.

Ludlow and Neale (1991) apply the same general considerations to the analysis of *indefinite* descriptions. According to Ludlow and Neale, there are a number of possible uses to which we can put indefinite descriptions, including referential uses, specific uses, definite uses, and purely existential uses. A *referential* use (as with definite descriptions) is a use in which an indefinite is used to refer to some particular individual (e.g., if I say 'A red-haired student in the front row is cheating on the exam' when there is one and only one red-haired student visible in the front row). A *specific* use is one in which I intend to communicate that I have singular grounds for my utterance but I don't intend my audience to identify the individual (as when, while proctoring an exam, I say 'A student is attempting to copy his neighbors' exams. Everyone please keep your work covered.'). A *definite* use is when I intend to communicate that exactly one individual satisfies the description, but I do not intend to communicate who the student is, or even that I know who it is ('A student robbed the answer sheet from my office leaving his fingerprints, but we haven't identified the prints'). A *purely existential*⁴ use is simply one that corresponds directly to the literal meaning of the existential quantifier ('The exam scores are too high. A student must have circulated the stolen answer sheet').

As argued by Ludlow and Neale, *all* of these uses can be derived from the simple quantificational interpretation of the indefinite. That is, semantically there is a

⁴ Ludlow and Neale use the phrase 'purely quantificational', but that phrase is more general than they intend.

unitary analysis of indefinite descriptions, but following the Gricean considerations highlighted by Kripke (1977) and Neale (1990), this single interpretation can be used in all of the ways just outlined.

Of particular interest to us now is the idea that pragmatics can (and routinely does) deliver both referential and definite uses of indefinite descriptions. What this means is that an indefinite description can be put to all the uses that a Russellian definite description can. But if that is so, then why do we need Russellian definite descriptions at all? Answer: we don't.

4. Extending the Gricean Strategy

As noted above, our central idea is that 'a' and 'the' are two expressions with different spellings but the same meanings—synonyms, rather like 'gray' and 'grizzled' or 'grisly' and 'gruesome'. Literally 'The F is G ' and 'An F is G ' both express the same thing as $[\exists x: Fx](Gx)$. When we use 'The F ' referentially, it is via the same mechanisms that we use $[\exists x: Fx]$ referentially. More significantly, we hold that even the uniqueness implication of 'The F ' is a matter of use, not what is literally said. The mechanisms that allow us to use 'The F ' definitely—i.e. to communicate that there is a unique satisfier of F —are similar to those that we employ in using 'An F ' definitely. In the next section, we will provide some detail about the nature of these mechanisms. First, however, we have to get clear on the true nature of the differences between 'a' and 'the'.

Obviously there are important differences between our application of the terms 'the' and 'a'. A student who did not know these differences would certainly have problems on his English language exams. But it does not follow that these differences in application are part of their semantical content. Indeed, many synonyms customarily *are* put to different uses.

Take, for example, Grice (1961, 1975) on the distinction between 'but' and 'and'. On Grice's story, 'but' and 'and' literally mean the same thing, but different 'conventional implicatures' are associated with them; 'but' implicates a sense of contrast between the conjuncts. Similarly it is reasonable to hold that 'unless' has the same semantics as 'or' in its exclusive use, but implicates that the proposition expressed by the initial conjunct is the more likely or default hypothesis. We want to tell a similar story about 'a' and 'the'. However, our story is *not* that the uniqueness implication of 'the' is part of its conventional implicature. We believe that the conventional implicature of 'the' is actually much weaker: it is simply that the noun-phrase with the determiner 'the' carries the conventional implication that the object under discussion is *given* in the conversational context. Noun-phrases fronted by the determiner 'a' are conventionally implicated to involve *new* information. The basic observation about new and given information is not novel, of course; many

382 *IV. Definites and Indefinites in Semantic Theory*

traditional studies of grammar make precisely this claim (see Christopherson 1939; Givon 1984), and the observation has not escaped the attention of lexicographers. Early on in the lengthy entry for ‘The’ in the *Shorter Oxford English Dictionary*, we find this entry: ‘the . . . Marking an object as before mentioned or already known or contextually particularized’.

Now clearly, just as when we use ‘but’ in a conjunction we are not literally expressing a claim that the first conjunct ordinarily contrasts with the second, we believe it is clear when we implicate that something is given information we are not explicitly saying: ‘this is given information’. It is simply something that competent users of English can infer from the conventional implicature inherent in ‘the’.⁵

Givenness is not the complete story, however. When the predicate by itself ensures uniqueness, the implicature gets overridden. If it is obvious that there can only be one *F*, then no issue can arise about singling out a specific *F* as already mentioned or contextually particularized. Consider, for example, superlatives (‘the tallest man in the room’) and titles (‘The President of the United States’).

There is no difference in truth conditions between ‘[ιx : *x* is the tallest man in the room](*Gx*)’ and ‘[$\exists x$: *x* tallest man in the room](*Gx*)’. The same point holds of titles. Given that there is, by definition, only one President of the United States, the choice of the logical determiner is utterly irrelevant to the truth conditions of the resulting sentence: ‘[ιx : *x* is the POTUS](*Gx*)’ has the same truth conditions as ‘[$\exists x$: *x* is the POTUS](*Gx*)’. In these cases, uniqueness is neither implicated, nor implied by the type of determiner. Rather it simply results from the descriptive material itself.⁶

5. The Formal Derivation

So far we have been describing in an informal way the Gricean mechanisms that give rise to uniqueness implications. We now want go into more detail about the

⁵ We are not wedded to the idea that this additional information be accounted for in terms of conventional implicature. We are aware that the notion of conventional implicature has fallen into disfavor in some quarters (see, e.g. Sperber and Wilson 1986/1995: 182; and Bach 1999*b*), and we note that all of what we say here can be recast in terms of explicature or inference or by one’s favorite account of how the new/given distinction is to be understood. Our key point is that the new/given distinction (however understood) plus pragmatic resources of some form will get us to definiteness. Anne Bezuidenhout has observed that it might be possible to build the new/given information into the *semantics* of the determiner. In principle we do not have a problem with that move either—it would mean that ‘the’ and ‘a’ *do* differ in meaning, but not in the way we ordinarily suppose.

⁶ Other cases where the givenness implications fail are the cases cited in examples (a)–(b) above. It is worth noting that many of these examples involve possessive notions of some form (the eye = his eye, the living-room wall = the living room’s wall, the bank = the river’s bank). One attractive proposal would be that sometimes the determiner signals inalienable possession of some form rather than givenness. As noted above, givenness is certainly not the whole story here.

nature of these mechanisms, and show how a uniqueness implication might be explicitly derived.

We will begin with an illustration, drawn from Neale (1990), of how such Gricean derivations work, and then we will proceed to give derivations for both the definite and referential uses of indefinite descriptions. All these derivations will fit the following template.

- (a) *S* has expressed the proposition that *p*.
- (b) There is no reason to suppose that *S* is not observing the Gricean cooperative principle (CP) and maxims of communication.
- (c) *S* could not be doing this unless he thought that *q*.
- (d) *S* knows (and knows that I know that he knows) that I can see that [he thinks]⁷ the supposition that he thinks that *q* is required.
- (e) *S* has done nothing to stop me thinking that *q*.
- (f) *S* intends me to think, or is at least willing to allow me to think, that *q*.
- (g) And so, *S* has implicated that *q*.

Neale's example of how this works involves a case where the speaker and hearer both know that Harry Smith is the chairman of the Flat Earth Society. The idea is that the speaker (*S*) can use the definite description 'The Chairman of the Flat Earth Society' to refer to Harry Smith. The derivation goes as follows:

- (a) *S* has expressed the proposition that $[\lambda x: Fx](Gx)$.
- (b) There is no reason to suppose that *S* is not observing the CP and maxims.
- (c) *S* could not be doing this unless he thought that *Gb* (where '*b*' is a name).

Gloss: On the assumption that *S* is observing the maxim of relation, he must be attempting to convey something beyond the general proposition that whoever is uniquely *F* is *G*. On the assumption that *S* is adhering to the maxim of quality, he must have adequate evidence for thinking that the *F* is *G*. I know *S* knows that *b* is the *F*; therefore *S* thinks that *Gb*.

- (d) *S* knows (and knows that I know that he knows) that I know that *b* is the *F*; that I know that *S* knows that *b* is the *F*; and that I can see that he thinks the supposition that he thinks that *Gb* is required.
- (e) *S* has done nothing to stop me thinking that *Gb*.
- (f) *S* intends me to think, or is at least willing to allow me to think, that *Gb*.
- (g) And so, *S* has implicated that *Gb*.

⁷ The additional 'he thinks' in square brackets is not in Grice's original formulation, but is added by Neale.

384 IV. *Definites and Indefinites in Semantic Theory*

In the case where we do not refer to an individual known by name, but rather to someone in the perceptual environment, we replace step *c* of the derivation with *c'*.

- (*c'*) *S* could not be doing this unless he thought that *Gb* (where '*b*' is a demonstrative that refers to an individual salient in the environment).

Gloss: On the assumption that *S* is observing the maxim of relation, he must be attempting to convey something beyond the general proposition that whoever is uniquely *F* is *G*. On the assumption that *S* is adhering to the maxim of quality, he must have adequate evidence for thinking that the *F* is *G*. It is not plausible to suppose that he has just general grounds for this belief, therefore he must have object-dependent grounds. I can see that there is someone in the perceptual environment who could be taken to satisfy the description 'the *F*', and I can see that *S* can see this. Therefore the grounds for his assertion that the *F* is *G* are plausibly furnished by the belief that *Gb*.

Crucially, we can also see that a similar derivation is possible for a *definite* use of an existentially quantified expression—that is, the use of an expression of the form $[\exists x: Fx](Gx)$ to communicate that there is a unique satisfier of the description. Let us return to the case in which *S* comes upon Smith's badly mutilated corpse and says 'The murderer of Smith is insane' and consider how the addressee might derive the implicature of uniqueness.

- (*a**) *S* has expressed the proposition that $[\exists x: Fx](Gx)$.
 (*b*) There is no reason to suppose that *S* is not observing the CP and maxims.
 (*c**) *S* could not be doing this unless he thought that $[\iota x: Fx](Gx)$.

Gloss: By invoking the determiner 'the', *S* intends to communicate that whatever *F* or *F*s he is talking about is/are given in the conversational context. By refraining from using a plural noun, *S* intends to communicate that just one *F* is given in the conversational context. If there were more than one *F* given in the context, *S* would have used the plural definite description (otherwise *S* would flout the maxim of quantity). But the only way a unique *F* could be given in the context, is if there is just one *F* in the domain of discourse as fixed by the current context. *S* knows this, and so must be assuming that there is exactly one *F* in the domain of discourse. Therefore the grounds for *S*'s literal assertion that $[\exists x: Fx](Gx)$ are plausibly furnished by the belief that $[\iota x: Fx](Gx)$.

- (*d**) *S* knows (and knows that I know that he knows) that I know that he thinks there is only one *F* and that I can see that *S* thinks the supposition that he thinks that $[\iota x: Fx](Gx)$ is required.
 (*e**) *S* has done nothing to stop me thinking that $[\iota x: Fx](Gx)$.
 (*f**) *S* intends me to think, or is at least willing to allow me to think, that $[\iota x: Fx](Gx)$.
 (*g**) And so, *S* has implicated that $[\iota x: Fx](Gx)$.

We conclude that there are no formal obstacles to the analysis of descriptions like ‘the *F*’ as mere existentially quantified noun-phrases. As we will see in the next section, there are good philosophical reasons for supposing that this is in fact the correct analysis.

6. Positive Evidence

If we pursue the path outlined above, then a number of familiar puzzles about definite descriptions begin to dissolve. These puzzles range from the problem of misdescription to the behavior of ‘the’ in conditionals to the general detachability of the uniqueness implications of ‘the’. In this section we will canvass some of these familiar problems with the standard Russellian view and show how an existential analysis of ‘the’ solves those problems.

The Argument from Misdescription

As Neale (1990: 91–3) observes, one of the great advantages of the Gricean distinction between the proposition said and the proposition meant, is that it offers an account for our being in two minds about certain famous ‘misdescription’ cases.

For example, in Donnellan’s case of Smith’s murderer discussed above, I might say ‘Smith’s murderer is insane’, and still say something true even if the crazy man at the defense table is entirely innocent of the charges and the actual murderer, who is miles from the scene, is quite sane. Of course, there is also some pull to say that in such a case one is saying something false as well. As Neale observes, this is simply a case where what we literally said was false, but what we intended to communicate—the proposition meant—was true. The two-level theory thus accounts for our conflicting intuitions.⁸

Neale observes that things work in the opposite direction as well. Hornsby (1977) gives the case of my observing the man ranting at the defense table and saying, ‘The murderer of Smith is insane’, not realizing that the man at the table is (i) innocent and (ii) quite sane, while the actual murderer is at large and certifiably insane. Again we are pulled in two directions about the truth of what I say, and again Neale observes that the distinction between the proposition meant and the proposition expressed allows us to explain why. In this case, the proposition expressed is literally true, but what I intended to communicate is mistaken.

The problem is that there are other cases of misdescription which remain problematic for the classical Russellian analysis of descriptions—in particular for attributive uses. Suppose, for example, that at the scene of the crime Detective

⁸ Subsequent work by Neale (1999) and a number of others has gone on to propose that two distinct propositions are being expressed. Everything we say here can be recast under that assumption.

386 *IV. Definites and Indefinites in Semantic Theory*

Watson says ‘The murderer of Smith is insane, whoever he is’. But now suppose that there is not one murderer, but two and that one or both of them is insane. Has Watson said something true or false? Again, we seem to be in two minds about this. There is a clear sense in which Watson is communicating something true, but there is also a sense in which his remarks are mistaken.

If we suppose that the English determiner ‘the’ actually has the semantic content of the existential quantifier, then our conflicting intuitions have a ready explanation. What Watson literally said was true; there is in fact a murderer of Smith who is insane. But what Watson was communicating to his audience—that there is only one such murderer—is mistaken.

The distinction between the proposition meant and the proposition expressed is a crucial part of the solution, but it cannot be the whole solution. Solving the problem of misdescription also requires that we give up the idea that uniqueness is part of the meaning of ‘the’. When Gricean mechanisms are combined with an existential analysis of the definite determiner, the residue of the argument from misdescription finally dissolves away.

Conditionals

Another class of constructions that has been a nagging problem for the classical Russellian analysis of ‘the’ has been its behavior in many conditional environments. The problem, very simply, is that the determiner fails to have the expected uniqueness implications in these environments. Consider the following examples from Ludlow (1994) and Moltmann (1996).⁹

- (1) If a doctor and a lawyer enter the room, the doctor will sit down first.
- (2) If anyone here has a dog, he has to register the dog.
- (3) If anyone has a dime, he should put the dime in the meter.
- (4) If a bishop meets another bishop, the bishop blesses the other bishop.

Obviously these sentences do not imply that there is a unique doctor, dog, dime, etc. What can account for this lack of uniqueness—particularly if uniqueness is supposed to be part of the meaning of the determiner? As a first stab, one might suppose that on a Lewis-type analysis of conditionals, the description will be relativized to a world, so that we will not be committed to the absurd view that these sentences entail that there is only one doctor, dog, dime, etc.

This move helps a little, but now the problem of incomplete descriptions returns with a vengeance.¹⁰ Since each nearby world is going to have several doctors, dogs,

⁹ The first, from Ludlow (1994) is attributed to Mark Johnson. The others, from Moltmann (1996) are attributed to Robert Fiengo. As Moltmann notes, similar phenomena also arise in so-called donkey-anaphora sentences. Consider: ‘Everyone who has a dog has to register the dog’.

¹⁰ The problem, as initially brought to attention by Strawson (1950) is that if I say ‘the table is covered with books’, I do not mean to be suggesting that there is only one table in the world, but that seems to be

dimes, etc., there is no unique doctor, dog, etc. to be picked out by the description. But the situation is even worse than standard solutions to the problem of incomplete descriptions—those which hold either (i) that the unique object can be picked out by fleshing out the description or (ii) that the unique object is picked out by restricting context.¹¹ Regarding strategy (ii), what sense does it make to talk of restricting the context in other worlds? Even if we *could* make sense of this idea, example (4) soon gives rise to new problems. What restriction on context is simultaneously going to include the event of a bishop blessing another bishop *and* allow that there is only one bishop in the specified context? This is like having a context in which we eat cucumber sandwiches, but in which there are no sandwiches.

Regarding strategy (i), there is simply no way to sufficiently flesh out the description, even employing the device of packing demonstratives into the description. Why? Because we crucially understand that any further descriptive facts concerning the two bishops, for example, are irrelevant to the truth of what we say. The bishops can be tall, short, Italian, or African, and they can certainly be well out of range of our demonstrative acts; it simply does not matter. Indeed it cannot matter, since in the case of conditionals we want our claim to cover all bishops—indeed, all possible bishops.

These considerations do not turn on the specific theory of conditionals that we choose to employ. For example, if we adopt an event-based analysis of conditionals in the spirit of Lycan (1984) and Kratzer (1989), we fare no better. (A version of this idea with application to descriptive pronouns is explored by Heim 1990*a*; Ludlow 1994; Moltmann 1996.) Applied to this case, the basic idea would be to

what the Russellian theory of descriptions is literally committed to. As Neale (1990) explains, two strategies have emerged for trying to answer this objection. The first strategy (adopted, e.g. by Sellars 1954) is that the unique object can be picked out by using context to flesh out the description. The second strategy is to say that the unique object is picked out by allowing the context to restrict the domain of quantification. See Lepore (this volume), Stanley and Szabó (2000*a*), Neale (2000) for a sample of the discussion.

¹¹ Both strategies face severe difficulties. The first one needs an account of what determines the features of context that complete description and how the hearer is supposed to discover them. For example, suppose that Smith enters a room, leaving the door open, and Brown comments: 'You left the door open'. There are indefinitely many possible completers for the description e.g. 'the door/that I am now staring at/through which you just entered/to this room/over there'. It is difficult to see what would make one of these, rather than any other, the real completer. Moreover, if Smith is to understand what Brown is saying, then he must know what the completer is. It is hard to see how he could know this. (for discussion see Wettstein 1981; Larson and Segal 1995.) The problem with the second strategy is that it requires an account of some rather special mechanisms of domain restriction. Consider the following example. We are watching a basketball game and a particular player, *P*, miraculously shakes off his defender, shoots, and scores. I say 'The man is a genius'. On the face of it, a number of men are in the domain of discourse. On the hypothesis that 'the' is a Russellian quantifier, the domain has shrunk to include just *P* and exclude all the other men. Compare the case in which we are watching a basketball game and several players cooperate in a masterful play. I cannot simply say 'Every man is a genius' and thereby felicitously communicate that every man who was involved in the masterful play is a genius. Nor can I say 'Most men are geniuses' to communicate e.g. the proposition that most men on the team that just scored are geniuses. These cases indicate that the proposed domain-shrinking mechanisms that account for incomplete descriptions are not the same as those that operate with other quantifiers, like 'most' and 'all'. Until some account of these special mechanisms is offered, the proposal remains merely promissory.

try and give an event/situation-based analysis of conditionals so that the definite descriptions with their uniqueness implications are always relativized to an event or situation. So e.g. in the case of (1) the gloss would be that ‘for every event in which a doctor and a lawyer enter the room there is a related event in which the unique doctor *in that event* will sit down first’. As noted by Heim and Ludlow, this solution gives rise to a number of puzzles that need to be confronted in the case of descriptive pronouns, and Moltmann argues that it fails altogether for both definite determiners and descriptive pronouns.¹²

Perhaps it is better to just cut the Gordian knot and reject outright the idea that the uniqueness clause is part of the meaning of ‘the’.

Possessives

Possessives, like ‘my *F*’ and ‘your *F*’ provide further evidence for the existential analysis of ‘the’. Since Russell (1905) these have typically been treated as equivalent to ‘The *F* of mine’, ‘the *F* of yours’, etc., and this treatment is vindicated by the explicit use of the definite article in possessive constructions in many languages: for example, Italian ‘Il mio libro’ (‘the my book’).

It seems to us, however, that these uses are more correctly analysed as being equivalent to $[\exists x: x \text{ is an } F \text{ of } \dots]$ (e.g. $[\exists x: x \text{ is a book of mine}]$). Obviously, if I say ‘This is my book’, or ‘She is my student’, or ‘You stole my idea’, there is not even the illusion of uniqueness. So why should we saddle our analysis of possessives with a clearly incorrect uniqueness clause?

It is of course possible to contend that this merely shows possessives to be indefinite descriptions and that it has nothing to do with the proper analysis of the English determiner ‘the’, but this contention strikes us as disingenuous. Languages with surface elements which are candidates for being genuine definite determiners are far and few between. In addition to English, the Romance languages offer some

¹² For example, one problem revolves around example (4). Very simply, no matter how finely we carve up events, we still need to be left with two bishops in the event in question. But then there simply cannot be a unique bishop in the event. (Heim attributes the observation of this problem to Hans Kamp and Jan van Eijk.) Ludlow proposes a possible solution that would allow the introduction of thematic roles into the actual descriptive content of the noun phrases, so that one would get ‘the bishop theta-1’ and ‘the bishop theta-2’. The only problem is that the usual theta-roles (agent, patient, theme) do not seem to be sufficient for distinguishing bishops that stand in the meeting relation (or any other symmetrical relation). One can introduce theta-role labels arbitrarily, but without some semantical content to those labels the proposal teeters on the brink of vacuousness. Another problem with the relativization-to-event strategy involves an extension of the ‘sage plant’ sentences discussed in Heim (1982) (Heim credits Barbara Partee with the original sage plant examples; Moltmann credits the observation about sage plant sentences with explicit descriptions to Bob Fiengo): ‘If a man bought a sage plant, he bought two others along with the sage plant.’ Suppose we take the description to be relativized to events, so we can gloss this as ‘for every event in which a man bought a sage plant, there is a related event in which he bought two others along with the sage plant *in that related event*’. The problem is that the consequent event has to contain at least three sage plants. Notice also that it will not help to say that ‘the sage plant’ is elliptical for ‘the sage plant he bought’ since by hypothesis there are *three* sage plants that he bought. Other maneuvers are possible (e.g. ‘the sage plant in question’ or the ‘sage plant in the antecedent event’) but how are the unique denotations of these descriptions to be identified? The problem is parallel to distinguishing the two bishops in (4).

candidates like the Italian 'il', 'la', 'lo' and the French 'le' and 'la'. The problem is that in Italian these elements appear explicitly in possessive constructions (as in 'il mio libro'—'the my book') yet carry no uniqueness implications. This fact undermines the plausibility of Italian determiners being genuine Russellian definite determiners. One can of course dismiss the Italian determiners and argue that the English 'the' is different, but then the defense of the Russellian nature of English 'the' looks like an ever lonelier task.

The 'Definiteness Effect'

At first blush the 'definiteness effect' appears to be a problem for the unitary analysis of descriptions. We concur with Szabó (2000) that it is not problematic, and we would go further and argue that, once properly understood, it provides more evidence for the unitary analysis.

The apparent problem is that definite determiners are prohibited from being inserted in certain grammatical environments, where indefinite determiners *are* permitted. The paradigm, first observed in Milsark (1974) involves so-called there-insertion contexts:

The basic observation from Milsark is that there-insertion is only possible with a particular class of determiners. So, for example, one gets the following distribution of facts.

- (5) There's a fox in the henhouse.
- (6) *There's the fox in the henhouse.
- (7) There are three foxes in the henhouse.
- (8) *There's every fox in the henhouse.
- (9) There are no foxes in the henhouse.
- (10) *There is every fox in the henhouse.

The crucial facts for our purposes concern (5) and (6). It appears that indefinite descriptions can be inserted in these environments and that definite descriptions cannot, and it is supposed that this is a reflex of something semantical.

A number of attempts have been made to give a semantical account of this phenomenon (see e.g. Barwise and Cooper 1981, Keenan and Stavi 1986, Higginbotham 1987) but none of these attempts have been entirely satisfactory.¹³ This failure should be no surprise if the considerations adduced above are correct. The effect cannot be semantical in character because there is no semantical difference between definite and indefinite descriptions.

There is, in fact, compelling evidence to support the contention that this is not a semantical effect. For example, a description with a uniqueness clause *can* be inserted in these environments.

- (11) There's exactly one fox in the hen house.

¹³ See Ludlow (1991) for discussion.

390 *IV. Definites and Indefinites in Semantic Theory*

This shows that it cannot be the uniqueness clause in ‘the’ that is blocking there-insertion.

One natural explanation for the distribution of there-insertion facts is that they are a reflection of pragmatics rather than semantics; perhaps (as suggested in McNally 1992, 1998; Zucci 1995; Ward and Birner 1995; and Lumsden 1988) it is the implication of givenness that is blocking there-insertion here (e.g. since ‘there is/are’ is a device typically used to implicate that new information is being introduced there is a kind of pragmatic clash when ‘there is’ signals new information and ‘the’ signals old information). This explanation is further supported by facts like the following:

- (12) *There are none of the foxes in the hen house.
 (13) There are none of the foxes in the universe of discourse in the hen-house.

In the first example ‘none of the foxes’ suggests that some foxes are already under discussion, which clashes with the implication of ‘there are’ that new information is being introduced. The second example shows a case where this givenness implication of the noun-phrase is defeated by the introduction of talk of the full universe of discourse, and hence, as predicted, there-insertion is possible despite our using the determiner ‘the’.

The upshot is that by unifying our semantical analysis of ‘the’ and ‘a’ we are able to see why attempts to give a semantical account of the differences in behavior between ‘the’ and ‘a’ consistently founder. It also steers us in the apparently more fruitful direction of investigating these facts as reflexes of pragmatics alone.¹⁴

Predicational Uses of Definite Descriptions

One of the central observations about indefinite descriptions is that they have predicational in addition to quantificational uses. For example, it has been observed that ‘a lawyer’ in the following example appears to be a simple predicate rather than a quantified noun-phrase.

- (14) John is a lawyer.

It has also been argued that ‘the lawyer’ does not have this sort of use.

- (15) #John is the lawyer.

¹⁴ But see Abbott (1999) for criticism of the givenness account of there-insertion contexts. Abbott’s positive proposal is that these sentences are blocked because the combination of an existence claim with the existential presupposition of the definite determiner creates a ‘conversational tension’—to wit: ‘it is not felicitous to assert something while presupposing it’. This story strikes us as untenable since one can perfectly well say ‘The mayor is a mayor’, and while it may be uninteresting it is not ungrammatical in the way that (6) is. Abbott’s account would also seem to falsely predict that (11) is ungrammatical, since ‘exactly one’ carries the same existential presupposition (or same existential entailment) that a definite description does.

So there is supposed to be an asymmetry between these two cases. If there is such an asymmetry, and if it is a *semantic* asymmetry, then it appears we have a reason for distinguishing 'the' and 'a' as being semantically distinct.

But is there a semantic asymmetry? We think not. In the first place, one can introduce a semantic equivalent of a Russellian description in a predicational environment and have a perfectly grammatical sentence.

(16) John is the unique satisfier of the predicate 'lawyer'.

As Smiley (1981) has observed, there are many cases in which descriptions of the form 'the *F*', when in object position, routinely behave like predicates.¹⁵

(17) John is becoming the biggest jerk in town.

(18) John was the mayor of Boston.

(19) John is the lawyer that I was talking about.

If one opts for a semantical asymmetry between definite and indefinite descriptions then these facts are difficult to explain.¹⁶ But if definite and indefinite descriptions differ pragmatically then we would expect there to be a pragmatic account of their different behaviour in predicate phrases. And such an account is not hard to find.

Earlier we noted that it is a conventional implicature of the use of 'the' is that it apply to given information. Anomalous cases, such as 'John is the lawyer' occur when the implicature of givenness is present but unfulfilled, that is, when no *F* (no lawyer) is given. Felicitous cases are those where either the implicature is cancelled (16, 17, and 18), or it is not cancelled but is fulfilled (19, 20, and 21):

(20) *a.* I heard that Smith offended a lawyer in court yesterday.

b. Yes, unfortunately John is the lawyer.

(21) *a.* I heard that Barry has joined your legal team as an expert in constitutional law.

b. No, Barry is the DNA expert.

Of course it might be argued that these cases with 'the' are not really cases of predication, but can be assimilated to identity statements:

(22) [the *x*: lawyer *x*] (John = *x*)

But notice that the same can be said for cases involving 'a lawyer'.

(23) [an *x*: lawyer *x*] (John = *x*)

Our interest here is not in defending the predicational theory of determiner phrases, but simply to point out that there is no interesting semantical asymmetry

¹⁵ For an interesting discussion and defense of the thesis that definite descriptions are always predicational, see Graff (2001).

¹⁶ See Enç (1986) for an attempt to explain these facts by positing implicit temporal indices within the descriptions (e.g. 'the Mayor of Boston' might be understood as 'the Mayor of Boston at time *t*'). This is certainly a possible gambit, but without serious changes it weds us to a semantics of tense and metaphysics of time that at least one author of this chapter would reject for independent reasons. See Ludlow (1999).

between the determiners as regards such environments. Or rather, that a number of philosophical puzzles can be avoided if we reject such asymmetries.

A Note on Focused Descriptions

Abbott (1999) has argued that certain focus facts involving definite descriptions tell against their analysis as being merely quantificational. Consider the following sentence.

(24) That wasn't *a* reason I left Pittsburgh, it was *the* reason.

According to Abbot, an utterance of (24), focused as indicated, implies that the speaker has only one reason for leaving Pittsburgh, not many. Now the idea is that this supports the uniqueness analysis of definite descriptions over the givenness analysis of definite descriptions. Stressing 'the' allows us to focus the contrast in meaning between 'the' and 'a' and our judgment is supposed to be that the contrast in these cases has to do with whether the description is unique or non-unique rather than whether it is given or novel. In effect, focus becomes a tool for probing the real difference in meaning in these cases and the evidence allegedly supports the thesis that the relevant difference in meaning is along the unique/non-unique scale.

Whatever might be said about the logic of this argument, we disagree with Abbott's judgment of what is happening in these cases. We take it that someone may very well utter (24) with stress as indicated despite having several reasons for leaving Pittsburgh. Stressing 'the' indicates that the reason in question was not merely one of many reasons for leaving but rather the causally determinate reason—the big reason.

In our view this is typically the case when the definite determiner is stressed. Consider the following example. Ludlow's third grade teacher had a husband named William Faulkner. On vacation in the South, he was asked, 'Are you *the* William Faulkner?' Presumably, the questioner was not asking if he was the unique individual named William Faulkner (at which point he was, since the author William Faulkner had long since passed on), but was asking whether this individual was the famous—i.e. given or familiar—William Faulkner.

A similar incident involves Ludlow's father, first name Robert. Being a sometime resident of Naples, Florida, he was often confused with the author Robert Ludlum (who lived in Naples and is only recently deceased). Persons would typically ask 'Are you *the* Robert Ludlow?' But presumably they were not asking whether he was the only Robert Ludlow in town (which he was), but whether he was identical to the famous—i.e. given or familiar—individual who was in fact named Robert Ludlum.

A final note: Abbott's paper includes a cartoon by Dick Guindon of the *Detroit Free Press* in which a precocious boy confronting a shopping mall Santa Claus says 'Not *the* Santa Claus'. Abbott seems to think the joke is that the boy is saying this is just one Santa out of many—he's merely *a* Santa Claus. We see the joke differently.

Isn't it actually the case that the boy is saying this isn't the real Santa Clause? That is, this isn't the Santa that he's read about in books and that is reputed to bring him toys? Not the given or familiar Santa?

7. Conclusion

We think the arguments marshaled in section 6 give strong support to the idea that 'the' and 'a' have the same literal meaning, yet a number of philosophers that we have spoken to find the very idea to be utterly crazy. But why? Ordinarily, analytic philosophers like to go where arguments take them, particularly if the logic of those arguments is not only familiar, but (as in this case) routinely employed.

'The' and 'a' are regarded differently, however. Russell's analyses of these constructions have been at the very center of analytic philosophy for nearly 100 years, and while many have disputed Russell's analysis of definite and indefinite descriptions, few have doubted that he was correct to give them *separate* analyses. This tradition is not enough to save their bifurcated treatment, however. As we saw in section 6, clinging doggedly to this tradition has led to legions of philosophical puzzles—puzzles that dissolve when the traditional view is abandoned.

Linguistic myopia cannot be discounted here. All too often we opt for the logical analysis that seems to fit the surface form of the language in which we work (i.e. English). But both Russell and Frege cautioned against this kind of slavishness to surface form. In this particular case, deference to surface form not only lands us in the midst of philosophical conundrums, but it lends a procrustean quality to our analyses when we try to extend them to the other languages of the world (in this case Slavic languages, Chinese, Japanese, etc.). If we are interested in the logical form of natural language (as opposed to the logical form of English only) we need analyses that 'travel well'. We cannot be satisfied that our analysis works for our own language if fails as an analysis of most other languages in the world. Accordingly, we urge that further consideration be paid to the unitary analysis.¹⁷

¹⁷ One topic we have not addressed here but which would certainly play a role in future investigations is the treatment of plural definite descriptions. If this analysis were extended to such constructions then the idea would be that a plural description like 'the dogs' would have the literal meaning of 'some dogs', but may be used to express something about all the dogs. *Prima facie* there is much to recommend such an analysis for plural descriptions. If I say that the dogs are barking tonight, I certainly do not mean that all of them are barking (and it is clear that this is not a generic usage—I may be talking only about dogs in the neighborhood). I mean that some dogs are barking. Less clear is what one would say about both (this example was brought to our attention by John Hawthorne). Clearly 'both dogs' does not seem to mean 'some of the two dogs'. Here we think there is a genuine numerical claim packed into the semantics ('both dogs are barking' must mean that two given dogs are barking) just as when we say 'the only dog is barking'. In this respect, English 'both' may pattern with the Italian expression for 'both'—'tutte e due' ('all and two'). Szabó (2000) considers this possibility and rejects it on the grounds that no single lexical item can do what 'all and two' can. Once again, our point is that these and similar cases deserve further study.

